

Appendix B

Selected Excerpts from

Review of the Revised Analytical Plan for EPA's Second Prospective Analysis – Benefits and Costs of the Clean Air Act 1990-2020, Draft Report, #EPA-SABCOUNCIL-ACV-XXX-XX, March 5, 2004.

Excerpt 1, from page 3 of the draft report: “Value of Premature Mortality and Morbidity Associated with Reductions in Air Pollution”

Excerpt 2, from pp 54- 62 of the draft report: “12. USE OF VSL META-ANALYSES”

Excerpt 3, from pp 121-127 of the draft report: “APPENDIX D: ADDITIONAL DISCUSSION CONCERNING THE USE OF VSLS”

1
2 Value of Premature Mortality and Morbidity Associated with Reductions in Air
3 Pollution: Uncertainty analysis with respect to Value of a Statistical Life (VSL) values
4 requires information about the distribution of VSL estimates corresponding to risks and
5 populations that are similar to those relevant for the CAAA. The marginal distribut ion of
6 all empirical VSL estimates derived across all contexts is unlikely to be appropriate for
7 this purpose, as is any arbitrary convenient assumption about distributional shape.

8
9 The Panel recommends a primary focus, at this juncture, on the Viscusi-Aldy
10 estimates of VSLs based on U.S. studies. The Agency should not rely exclusively on the
11 Kochi et al. meta-analysis, which has not yet been peer-reviewed and published.

12
13 The Council Special Panel does not support an effort by the Agency to comply
14 with the OMB requirement for cost-effectiveness analysis by utilizing Quality-Adjusted
15 Life Year (QALY) as the measure of effectiveness. Too many other classes of benefits
16 besides human health benefits must be taken into consideration. A workshop on
17 appropriate cost-effectiveness approaches for this application may be helpful, but its
18 scope would need to be very carefully defined and the differences between cost-
19 effectiveness analysis in the typical health context versus cost-effectiveness for specific
20 human health benefits of the Clean Air Act (CAA) would be an important dimension of
21 the discussion.

22
23 Concerning morbidity, the Agency should continue to use Willingness-To-Pay
24 (WTP) estimates for morbidity values, rather than COI estimates, should these be
25 available. Where WTP is unavailable, COI estimates can be used as placeholders,
26 awaiting further research, provided these decisions include suitable caveats. The Dickie
27 and Ulery study is a valuable addition to the repertoire of empirical results concerning
28 WTP for acute respiratory illnesses and symptoms, although it is not so superior as to
29 supercede all earlier studies.

30
31 Ecological Effects: Human health risk reductions may be the most substantial
32 benefit from the CAAA, but they are not the only important benefit. Benefits to
33 ecosystems and other welfare benefits such as visibility are likely to be substantial and
34 are still receiving limited attention. The Council nevertheless recognizes substantial
35 challenges in quantitative assessment of these benefits. The greater heterogeneity in
36 ecosystems services makes it even more difficult to produce estimates of the benefits
37 from their protection than for the protection of human health. The input of the new
38 Science Advisory Board Committee on Valuing the Protection of Ecological Systems and
39 Services (CVPESS) and a new Council Ecological Effects Subcommittee (EES) may be
40 able to stimulate the development of greater expertise on this issue than is presently
41 available. Ecological effects to be valued must be limited to those effects for which there
42 is a defensible, rather than just speculative, link between air emissions and service flows.
43 The Council strongly objects to using inappropriate or unsupported placeholder values in
44 the absence of better information.

12.USE OF VSL META-ANALYSES

12.1. Agency Charge Questions Related to Use of VSL Meta-Analysis.

Charge Question 22: EPA's current analytic blueprint calls for an expert-judgment project on VSL determination that would produce a probability distribution over the range of possible VSL values for use in the 812 project. EPA is not sure how much priority to give to this project. A much simpler alternative would be for EPA to specify a plausible range of VSL values. One option would be to use a range bounded by \$1 million (based roughly on the lower bound of the interquartile range from the Mrozek-Taylor meta-analysis) and \$10 million (based roughly on the upper bound of the interquartile range of the Viscusi-Aldy meta-analysis. This range would match that reflected in EPA's sensitivity analysis of the alternative benefit estimate for the off-road diesel rulemaking. The range would then be characterized using a normal, half-cosine, uniform or triangular distribution over that range of VSL values. EPA would then ask this Committee to review this distribution. This approach could be done relatively quickly, based on the reviews and meta-analyses commissioned to date, and would allow a formal probability analysis to proceed, without suggesting that the Agency is trying to bring more precision to this issue than is warranted by the available science.

Charge Question 23: Pursuant to SAB Council advice from the review of the first draft analytical blueprint, EPA reviewed a number of meta-analyses –either completed or underway– developed to provide estimates for the value of statistical life (VSL) to be applied in the current study. EPA plans to consult with the Council (and coordinate this consultation with the EEAC) on how best to incorporate information from the Kochi et al (2002) meta-analysis, other published meta-analyses (Mrozek and Taylor and Viscusi and Aldy), and recent published research to develop estimates of VSL for use in this study. In addition, EPA plans to implement two particular adjustments to the core VSL values: discounting of lagged effects and longitudinal adjustment to reflect changes in aggregate income. Does the Council support these plans, including the specific plans for the adjustments described in chapter 8? If the Council does not support these plans, are there alternative data or methods the Council recommends?

Charge Question 31: EPA plans to work with the Council and the EEAC to develop revised guidance on appropriate VSL measures. We hope to include the Kochi et al (2002) meta-analysis, other recent meta-analysis, recent publications, and the 3 literature reviews sponsored by EPA. (A separate charge question pertaining to this element of EPA's VSL plan is presented below). In addition, EPA plans to conduct a follow-on meta-regression analysis of the existing VSL literature to provide insight into the systematic impacts of study design attributes, risk characteristics, and population attributes on the mean and variance of VSL. Does the Council support the plans described in chapter 9 for conducting this meta-regression analysis? If the Council does not support this analysis or any particular aspect of its design, are there alternative approaches which the Council recommends for quantifying the impact of study design

1 attributes, risk characteristics, and population attributes on the mean and variance of
2 VSL?

3
4 Charge Question 37: Does the Council support including the Kochi et al. (2002)
5 meta-analysis as part of a larger data base of studies to derive an estimate for the value of
6 avoided premature mortality attributable to air pollution? Are there additional data,
7 models, or studies the Council recommends? Does the SAB think that EPA should
8 include Kochi et al. 2003 if not accepted for publication in a peer reviewed journal by the
9 time the final 812 report is completed?

10 **12.2. Summary of Council Response**

11 The Council has combined the responses to charge questions 22, 23, 31, and 37
12 and has provided additional discussion concerning the use of VSLs in Appendix B of this
13 Council Report. Major summary points appear below.

- 14
15 - Since the Panel's initial receipt of the Analytical Plan, the plan for an expert-
16 judgment project on VSLs has been dropped from the blueprint. The expert
17 elicitation exercise is no longer an active portion of this charge question.
18
- 19 - Uncertainty analysis with respect to VSL values requires information about
20 the distribution of VSL estimates corresponding to risks and populations that
21 are similar to those relevant for the CAAA. The univariate distribution of all
22 empirical VSL point estimates derived across all contexts is unlikely to be
23 appropriate for this purpose, as is any arbitrary convenient distributional
24 shape.
25
- 26 - Discounting of lagged effects is advisable, but the literature on discount rates
27 for future financial outcomes and future health states is not clear on whether
28 straightforward discounting using an exponential model and a common rate
29 will be appropriate. Sensitivity analysis and caveats are recommended.
30
- 31 - Adjustments for future changes in aggregate income levels are being based on
32 very limited empirical evidence and should be considered placeholder efforts
33 at present. It would be preferable in the future if these adjustments were made
34 in the context of a formal model of preferences and the relevant elasticities.
35 Placeholder efforts should be clearly identified as such, and accompanied by
36 strong caveats. The First Prospective Analysis included (in an Appendix)
37 estimates allowing income growth. This type of analysis may be a candidate
38 for the recommended "exploratory" or preliminary analyses discussed earlier.
39
- 40 - The Panel recommends a primary focus, at this juncture, on the Viscusi-Aldy
41 estimates based on U.S. studies, although work in the direction of the Kochi et
42 al. analysis should be encouraged. Preferably, the variance estimates should
43 be based on the variance in the conditional expectation from the model, for a
44 set of conditions that most closely approximate those relevant for the CAAA.

- It is certainly reasonable to expect that the Second Prospective Analysis would consider insights derived from the other VSL meta-analyses (e.g. Mrozek and Taylor, and Kochi et al.). The Council recommends that, to the extent VSL measures are developed as conditional expectations from a meta-analysis, they should rely primarily on published peer review studies. As the Council's general comments on approaches to methodological innovation imply, the meta-analyses that best serve Agency needs will not always be published.
- Continual evolution of the relevant literatures justifies development by the Agency of a more formal laboratory phase for evaluation of potential methodological innovations. A "satellite benefit-cost analysis" based on updated methodologies could serve as a forum for evaluation of new methods before these innovations are formally and widely adopted by the Agency for the Section 812 Analyses and other analyses.

12.3. Expert Judgment - VSLs

The Agency desires to bound the range of plausible VSL values between \$1 million and \$10 million, which seems reasonable given the state of knowledge about empirical values in different contexts. This range, however, represents the marginal distribution of VSL estimates aggregated across values that have been determined in very different contexts. The ideal VSL distribution to employ would be the conditional distribution of VSL values, derived for contexts that most closely match the risks and affected populations relevant to the CAAA. This VSL does not necessarily lie in the middle of the overall marginal distribution of empirical VSL estimates across the broad range of contexts in the literature.

Some VSL distribution is needed from which to draw alternative point values of the VSL for simulations of the effect of uncertainty about VSL values. However, the Council Special Panel does not agree with arbitrary assignment of some convenient distribution (e.g. normal, half-cosine, uniform or triangular) for the range of values. Why not compare Mrozek-Taylor versus Viscusi-Aldy meta-analyses, including the latter's re-estimates with a sample consisting of one observation per study? Use these estimates to derive an appropriate mean and variance of the relevant conditional distribution from that model "configured" for the policy analysis. The idea is to narrow the range of plausible VSL estimates to reflect more closely the risks and affected populations for the policies in question.

12.4. Adjusting for latencies, income growth?

Latency in health effects, as well as cessation lags, mean that a comprehensive assessment of mortality risk reduction benefits must take into account individual discounting. In discounting individual health effects, there remains an important question as to whether the usual convenient exponential form of discounting is an appropriate assumption, given the numerous empirical anomalies. There are also unresolved

1 questions about the difference in discount rates concerning future health, as opposed to
2 future financial status. While the Council concurs that future benefits need to be
3 discounted, there is no consensus in the literature concerning how to do this. As a
4 practical matter, pending additional research, the Agency should adopt discounting
5 assumptions that are consistent with the rest of the Analytical Plan and include sensitivity
6 analysis and caveats.

7
8 The Panel does not support the use of the proposed adjustment for aggregate
9 income growth. This is arbitrary and inconsistent with VSL as a marginal rate of
10 substitution (MRS). The Council acknowledges that, in principle, demands for
11 environmental risk reductions (like demands for all other goods and services) are likely to
12 vary systematically across individuals with such factors as income, age, gender, ethnicity,
13 or a host of other variables. However, empirical evidence based upon utility-theoretic
14 specifications has not yet been amassed to a point where there is any professional
15 consensus as to the precise way in which demand for risk reductions varies with these
16 factors. The Council also acknowledges methodological change without full vetting and
17 review runs the risk of creating an appearance of manipulation. Thus, it is imperative
18 that the Agency substantiate any adjustments before attempting to incorporate them in the
19 Section 812 Analyses.

20
21 The Agency needs to be aware that there are some important subtleties concerning
22 income in revealed preference derivations of the marginal rate of substitution between
23 risk reductions and income. Income adjustments to VSLs (or equivalently to marginal
24 rates of substitution) require very stringent approximations. While empirical evidence for
25 income effects is substantial, it is generally derived from *ad hoc* reduced-form
26 specifications, rather than any formal theoretical basis.

27
28 Nonetheless, it remains clear that the Agency should take into account that, over
29 time, average real incomes are likely to grow. The Agency should continue to consider
30 ways in which to capture overall real income growth. Unfortunately, most of the
31 literature on income elasticities in VSLs is not based upon a framework that produces
32 reliable estimates of what adjustments should be made in the aggregate, over time. The
33 Council cannot support the proposed adjustments for aggregate income growth as being
34 theoretically consistent.

35
36 Any income adjustments in the present analysis fall within the category of
37 satellite or exploratory analyses that may be developed as supplementary to the primary
38 analysis. As such, they would be intended to stimulate discussion and review, rather than
39 constituting a primary component of an analysis intended to be used in evaluating a
40 policy. In any provisional analysis, it may be possible to place bounds on the likely
41 errors that would accompany simple approximations to likely income effects. If an
42 adjustment of this type is considered essential even at this stage in the analytical process,
43 the Agency should be especially prudent in qualifying it and present the results in a
44 format that is as transparent as possible. This would include explaining in detail how any
45 income adjustments have been accomplished and why they are deemed to be necessary.

1 It is worth emphasizing that as soon as the Agency begins to manipulate VSL
2 estimates to reflect anticipated changes in real incomes, it opens the door to arguments
3 that VSLs should also be adjusted for other long-run changes. These might include other
4 changes in budget constraints, such as alterations to the relative prices of medical care.
5 Or, they could include shifts in typical indicators of preferences, such as trends in the
6 sociodemographic mix in the population (e.g. changes in the age distribution).

8 The Agency should also be aware that if VSLs are to be adjusted for income
9 growth, so should be all of the other demand-based benefit measurements entertained in
10 the Section 812 Analyses. It may be difficult to defend making income-growth
11 adjustments only to one component on the benefits algebra.

13 In the longer term, consideration should be given to obtaining income-based
14 adjustments to VSLs (or even other types of adjustments) through preference calibration
15 techniques. These methods hold promise for generating forecasts that are consistent with
16 the relevant elasticities (see Smith, Pattanayak, and Van Houtven, 2003).

17 **12.5. Available meta-analyses**

18 Three meta-analyses were discussed in EPA's evaluation of summary measures
19 for the available VSL estimates (Mrozek and Taylor, 2002, Viscusi and Aldy, 2003, and
20 Kochi, Hubbell, and Kramer, 2003). The studies differ in several key respects, including:

- 22 a. The number of observations included from each study;
- 23 b. The format of the observations (e.g. actual estimates, use of group means,
24 and other transformations of the primary estimates);
- 25 c. The sample composition – U.S. studies, international, revealed and stated
26 preference;
- 27 d. The set of independent variables used for controls (e.g. inclusion of
28 industry effects);
- 29 e. Bayesian means versus regression summaries;
- 30 f. Published versus unpublished summaries.

32 The background for the charge questions tends to focus attention on the selection
33 of a single study as a summary for developing for the Prospective Analysis “one” VSL
34 estimate of reductions in mortality risk. However, the charge questions explicitly refer to
35 the “systematic impacts of study design attributes, risk characteristics, and population
36 attributes on the mean and variance of VSL.” The earlier meta-analysis strategies tended
37 to miss the opportunity to combine the insights from all studies to influence how
38 summary measures are constructed and used. We recommend that serious consideration
39 be given to using these insights in adapting how any meta-summary is used.

41 Equally important, the sensitivity of VSL estimates from meta-summary
42 equations to the sample composition (i.e. which studies are included) and to the controls
43 used (i.e. which study features are explicitly modeled) suggests that it would be prudent
44 to use the resulting lessons from this research in at least three ways:

- a. If one study, such as the Viscusi and Aldy (2003) meta-analysis, is selected, evaluate the sensitivity of the conditional expectation to the baseline risk and other control variables selected in measuring the conditional prediction.
- b. Evaluate the variance in the conditional prediction as a function of the values for the independent variables included in the model in relation to the mean values for these variables for the sample used to estimate the model.
- c. Consider the effects of inclusion or exclusion of independent variables or observations on the coefficient estimate for the risk measure. The data sets used in these studies are generally available for attempts at replication, so this type of comparison can be readily undertaken and would permit evaluation of the sensitivity of the VSL estimate to assumptions made, based on the available literature.

In general, it does not seem prudent to extend the sample to include studies for labor markets outside the U.S. The terms of employment, information about safety conditions, fringe benefits (e.g. health insurance), etc. are likely to be so different that one could not be sure that differences attributed to income or risk levels were in fact due to these variables.

12.6. Interpreting CV measures as opposed to wage-risk measures

One advantage asserted for the Kochi et al. study is the inclusion of contingent valuation (CV) evidence concerning VSLs. However, there is an important issue that has not been adequately discussed when CV results are included with revealed-preference wage-risk results concerning VSLs. The CV based measure of the VSL implicitly accepts a proportionality assumption between ex ante willingness to pay and the risk change.

The proper theoretical interpretation of the CV measures is as an *ex ante* option price for a risk change. If OP denotes the value for a risk reduction from P0 to P1 (with $P1 < P0$), and the P's designate the probability of death before and after the risk reduction, theory implies:

$$OP = f(P0, P1, \text{ and other variables})$$

The comma between P0 and P1 implies that linear proportionality in (P0 - P1) is an approximation, not a feature implied by theory. Thus, to rewrite equation (1) as equation (2) below, where the option price associated with a risk reduction is proportional to the size of the risk reduction (as well as being a function of a number of other variables) and then to approximate VSL as in equation (3) by normalizing upon a 1.00 risk change, adds additional untested assumptions.

$$OP = (P_0 - P_1) \cdot g(\text{other variables})$$

$$VSL \approx \frac{OP}{(P_0 - P_1)} = g(\text{other variables})$$

A meta-analysis that includes CV studies to expand the range of risk changes (or the types of risks considered) will accomplish this objective. However, it also changes the summary measure from an *ex ante* marginal rate of substitution to a linear approximation. Unfortunately, this added condition makes it difficult to evaluate whether the resulting differences in summary results between CV and wage-risk studies should be attributed to these additional assumptions implicitly added to the model, or to the expansion in the range or types of risks.

Nevertheless, the Council recognizes that CV-based studies offer unique opportunities to examine the empirical influence of many additional factors on the resulting estimates of VSLs. Despite the potential difficulty in rendering their findings compatible with those from revealed-preference wage-risk studies, CV studies have the potential to make important contributions to our understanding of how consumers value risk reductions, and it is important to take advantage of these opportunities.

12.7. Emerging considerations

As recent unpublished research by Cameron and DeShazo seems to suggest, the terms identified in equations (1), (2), and (3) above, and other things, may well be very important to the *ex ante* option price measured for the risk change. This research is presently available only as early reports from a detailed contingent valuation study. Nonetheless, it reaffirms the notion that it may be important to evaluate the sensitivity of the conditional expectation of the VSL to the conditioning variables used in its construction.

The Council's discussion also supported efforts to refocus attention on incremental willingness to pay for an incremental risk change, rather than the traditional, but potentially confusing construct that is a VSL. The panel's discussion urged EPA to consider including a preamble on the concept that is sought as a benefit measure, its likely link to the conditions of daily living and illness preceding death, as well as to any latency and temporal issues associated with exposure and increased risk of death.

The Panel recognizes that the current state of research makes it unlikely that empirical measures can imminently be developed that reflect all of these concerns. Nonetheless, the discussion led to a consensus that the Panel should urge Agency staff to consider careful qualification and sensitivity analysis for the measure used to monetize mortality risk reductions.

12.8. Which meta-analyses to use

In general, the Council Special Panel recommends that the Kochi et al. meta-analysis should not be given any particular prominence among the alternative meta-analyses used for determining one appropriate measure to use for the VSL. There are several reasons:

- a. The Kochi study is still unpublished. While it can sometimes be difficult to publish further meta-analyses when others are already in the literature, the Agency should not rely disproportionately on the Kochi study before it has been thoroughly peer-reviewed. The standards for peer-review obviously differ across journals and even across reviewers, but reliable peer-review can also be accomplished outside of the journal publication process. Both Mrozek and Taylor (2001) and Viscusi and Aldy (2003), however, have already appeared in the peer-reviewed literature.
- b. There are problems in the derivation of the variance of the VSL estimates. Some appear to be typographical errors. The researchers apparently faced some problems in terms of unobserved (or unreported) covariances among parameter estimates. However, it might be possible to derive estimates of variance in mean annual wage from the current population survey (CPS) or other sources, and use this information to fill in some of the blanks. It is not clear whether one should use a predicted wage or an actual mean wage. Overall, this is a careful study but, like all meta-analyses, it needs to address the potential impact of some of its key assumptions on the results of the analysis before it is possible to assess their importance.
- c. The use of author-specific means of VSL (p. H-12 to H-13) is troublesome if the different estimates have been derived from different samples.

If called upon to recommend just a single meta-analysis at this point, the Council Panel would recommend a primary focus on the Viscusi-Aldy estimates based on U.S. studies. However, as the 812 process evolves over time, the Council has recommended a commitment to Satellite or provisional analysis to test new methods in a policy relevant format. This would assure that the Agency did not miss opportunities to incorporate insights from new research as it emerges. It would also signal a commitment to understanding the full implications of methodology change before it was adopted as the “Agency Practice.”

Finally, variance estimates for the VSL measures predicted for a risk context and an affected population similar to those relevant to the CAAA should be based on the variance in the conditional expectation from the model.

12.9. Unpublished meta-analyses?

The Council was asked explicitly to address the question of unpublished meta-analyses. In general, we believe a peer-reviewed study will have greater professional credibility than one that has not met this standard. The Panel has some reservations about basing an analysis with the gravity of the Second Prospective Analysis on unpublished research, but has even greater reservations about using entirely non-peer-reviewed research. Each of the available meta-analytic studies has different advantages and shortcomings so that no single study should be the sole basis for information about the distribution to be used for the VSL in the Second Prospective Analysis.

This is another reason for creating an ongoing commitment by the Agency to engage in activities that serve as laboratories for methodological developments. Based on innovations in the literature, new methods and new meta-analyses will continue to be developed and applied to policy issues. First, they should be used for evaluative purposes. Results designated as explicitly as “exploratory” can be disseminated in Agency working papers to evaluate the implications of new proposals for analysis. This process serves a role that parallels the peer review process. However, it is more focused and relevant to Agency needs because the appropriate policy context is being considered. These satellite benefit cost analyses could then provide a forum for exchange and evaluation of new methods before they are formally adopted for specific analyses that would be submitted as the Agency’s official evaluation of a proposed regulation.

APPENDIX D: ADDITIONAL DISCUSSION CONCERNING THE USE OF VSLs

This appendix covers material that can be classified as “experimental” or “methods development.” It emphasizes some shortcomings of existing practices with respect to VSLs. The Agency is advised to anticipate changes in the state of the art in human health benefits valuation that may be appropriate to incorporate in future analyses as these updated approaches are vetted and as the justification for them becomes more widely understood.

The Council first wishes to highlight persistent conceptual problems stemming from the use of “the VSL.” Normalizing WTP to a 1.00 risk reduction is arbitrary and has proven to be confusing to non-specialists and therefore open to being used in a strategically misleading fashion. As a device for combining WTP estimates based on different risk changes, any arbitrary normalization is equally appropriate and a more policy-relevant risk change would be preferable for normalization, even if this necessitates a change in traditions.

That WTP should be close to proportional to the size of the risk change has theoretical support and would be enormously convenient. However, empirical tests of this theory are very difficult with hedonic wage data and contingent valuation studies tend to produce results at odds with this assumption. More information on this important aspect of VSL implementation would be valuable.

WTP for risk reductions should be presumed to be heterogeneous across risks and individuals, unless demonstrated otherwise. It is important that the proposed meta-analyses are designed to recognize this.

Existing meta-analyses have tended to maintain the hypothesis that there exists a single immutable VSL (or a simple VSL function that depends mostly on income levels). The early Agency posture suggested that this unknown VSL merely needed to be revealed by somehow combining VSL estimates from different studies.

The studies that form the raw material for meta-analysis may be compromised to varying degrees by their subjects having had incomplete information about risk. Credible meta-analyses should address these problems as well.

The Agency should proceed cautiously in adopting the results of existing or new meta-analyses as the basis for some assumed distribution for the WTP that will be appropriate for the Second Prospective Analysis. The contexts of the constituent studies may not adequately match the policy context where the WTP is needed.

D.1. VSLs vs. Micromorts

The concept of the value of a statistical life has unnecessarily impeded clear communication with risk managers about the public's value for small changes in health risks. However, the Council acknowledges that it is not in the Agency's best interest to attempt to take the lead by proposing fundamental changes in the way economists traditionally have thought about valuing mortality risks. Such initiatives properly comes from the academic community. However, the Council wishes to draw the Agency's attention to ideas and approaches that are likely to develop in the literature over the next few years. Even without adopting a substantially different perspective on mortality risk valuation, the Agency can report mortality values in ways that are less susceptible to misinterpretation by non-experts in the constituency for the Section 812 reports. Specifically, the Agency should exercise more precision in describing and qualifying the measures of mortality risk reduction it currently uses. Whenever the concept of a VSL is introduced, the Agency should identify the VSL explicitly as a normalization relative to a particular baseline risk. The corresponding range of untransformed WTP estimates for the policy-relevant range of risk changes should be provided for comparison.

VSL is defined as the marginal rate of substitution (MRS), namely the (local) difference in income that will leave an individual equally well off in the face of a difference in mortality risk. It is well recognized in the literature that this MRS depends on baseline risk, income, and may well depend on other characteristics of the risk and the individual. The units in which this MRS is described are arbitrary (e.g., dollars per pound, pennies per ton, etc.). By focusing on "the Value of a Statistical Life," we have arbitrarily adopted as our units "dollars per 1.00 risk change."


The population WTP for a specified risk reduction is defined as the sum of individuals' WTP for the individual risk reductions. For example, if a policy change reduces fatality risk this year by Δr for everyone in a population of size N , the population WTP for this change can be calculated as vN , where v is the population average WTP for a Δr reduction in the chance of dying this year. This same population value is often described as the product of the average VSL and the expected number of "lives saved" by the risk reduction. Using the normalization of dollars per 1.0 risk change, VSL is defined as $v / \Delta r$, and "lives saved" is equal to the expected number of deaths averted this year, i.e., $N \Delta r$.

While this alternative formulation, in terms of the average VSL and the number of "lives saved," is mathematically equivalent to the population WTP (i.e., the product of the average WTP and the population size), it is potentially misleading. It suggests that the value of each "life saved" is equal to the average VSL, and that one only needs to know the expected number of "lives saved" in order to calculate population WTP. In addition to other factors, VSL is likely to depend on the size of the individual risk reduction Δr , and so the population WTP for a change that "saves one life" may depend on whether the change reduces many people's risk by a small amount or reduces a small number of people's risk by a large amount.

The arbitrary choices made with respect to the normalization of VSLs unnecessarily court objections from non-specialists who confuse “The Value of a Statistical Life” (the economists’ technical term for an extrapolated linear approximation to a marginal measure) with “The Value of Life” in the sense of some measure of the intrinsic value of one human life with certainty. Long ago, Ron Howard (1984) proposed the term “micromort,” meaning the value of a one-in-a-million risk reduction, which would translate into one one-millionth of our usual \$5-6 million VSL, or just 5 to 6 dollars. This metric would be less misleading than the VSL, but unfortunately it has never achieved currency. There is no imperative to choose a 1.00 risk change as the intervening metric for scaling. Scaling all estimates to the risk change relevant for some specific policy is just as valid, and would lead to the identical mathematical result for aggregate WTP for a risk reduction policy.

There are other potential concerns about empirical measures of WTP for risk reductions. Suppose that we are trying to combine the information about WTP for risk reductions from five different studies, each involving one particular (different) risk reduction, r_1 through r_5 , as in the figure. (With any luck, there will be standard errors on the underlying WTP estimates, as shown, so there will be corresponding standard errors on the resulting individual studies’ estimates of VSLs, although these are not depicted in the diagram.)

If we use the WTP and risk information from each study to impute the associated VSL for a 1.00 risk change, the numbers may vary widely, as shown. It is these different VSL estimates that most meta-analyses seek to “average” according to formulas of different complexity and sophistication. By taking some type of average of the five separate VSLs, we can infer an average WTP for risk reductions that controls for the different risks across studies. However, if the true WTP function tracks along the dashed line, and if the policy context concerns a risk change that is, say, slightly larger than r_5 , then the WTP that would be inferred from the average VSL would be an inappropriate estimate.



1 The individual WTP point values depicted in the diagram may also differ because
2 of other types of heterogeneity across the contexts wherein they were derived. In that
3 case, it would of course be inappropriate to average these results, even after
4 normalization to a common 1.00 risk change.

5
6 VSLs are based on empirical data concerning choices in the neighborhood of very
7 small risks and small risk differences. Outside of this domain, we can really say nothing
8 about WTP for much larger risks and risk changes. The implicit extrapolation to a 1.00
9 risk change that produces a VSL is understood by specialists to be purely a convenient
10 device to control for variations in the sizes of risk reductions across the studies that yield
11 these estimates. Unfortunately, this is often not understood as such by non-specialists.
12

13 **D.2. Proportionality**

14
15 The VSL can be viewed simply as a strategy for getting around the fact that WTP
16 from different studies corresponds to different sized risk changes. It would be
17 inappropriate to average the individual WTP estimates without acknowledging that they
18 apply to different risk changes. The issue of proportionality of estimated WTP for risk
19 reduction and magnitudes of these risk reductions has been raised previously (e.g.
20 Hammitt and Graham, 1999). Certainly, if we wish to maintain the hypothesis that there
21 exists a single one-size-fits-all VSL that is the same for all possible risk reductions, then
22 the estimated WTP for different risk reductions ought to be proportional to the sizes of
23 the risk reductions in question. This constitutes a requirement for a very specific type of
24 “scope test.” However, not all empirical estimates of WTP functions produce parameters
25 that are consistent with this requirement. Some studies show negligible effects of risk
26 changes on WTP. Such a result that is clearly problematic for valuing mortality risks.
27 However, other studies reveal estimates that suggest that WTP is not strictly proportional
28 to the size of the risk change.
29

30 Stated-preference (e.g. contingent valuation) studies almost invariably show that
31 WTP is an increasing but concave function of risk reduction. Revealed-preference studies
32 (e.g., hedonic wage studies) typically do not tell us anything about how WTP depends on
33 the magnitude of the risk change because we model workers as choosing jobs from a
34 continuous set of jobs that differ in wage and risk, and typically do not have information
35 on what jobs (and risks) and individual rejects.
36

37 For example, compensating-wage-differential estimates are based on fitting a
38 regression model to data on individual workers’ wages, occupational fatality risks, and
39 other variables such as education and job experience that influence wages. This
40 regression estimates how wages vary with occupational fatality risk, holding other factors
41 constant. Each worker is assumed to prefer the job he holds to other jobs that are
42 potentially available to him, which are characterized by the regression. Setting the

independent variables equal to the worker's characteristics, the regression is interpreted as describing how the set of jobs available to him differ in wage and risk.

Many of the studies that yield WTP estimates do so for only a single common risk difference for all subjects, so there is too little information in any single study to assess the effect of the size of the risk change on WTP. Some sort of preference calibration exercise would be necessary in order to combine all of the available estimates.

D.3. Heterogeneity: Context-dependent WTP

Many practitioners seem to lose sight of the subtlety that the VSL is not a physical constant, like the constant of gravitation $(6.673 \pm 0.003) \times 10^{-8} \text{ cm}^3\text{gm}^{-1}\text{s}^{-2}$, or the mass of a hydrogen atom $(1.67339 \pm 0.0031) \times 10^{-24} \text{ g}$. Instead, VSL is an artifact of human preferences. It is based on willingness to pay for risk reduction, which depends on the marginal (dis)utility of risk and on the marginal utility of income. While it may be possible to identify some regularities across types of people in these two marginal utilities, it is conceivable that they are essentially unique to each person. Therefore, so can be the corresponding VSL.

The contexts for empirical studies concerning risk tradeoffs differ in many more ways besides just the risk change they consider. The types of risk and the characteristics of the individuals experiencing these risks can also lead to heterogeneity in WTP. If the policy context is not "in the middle" of the range of study contexts, then it can be potentially very misleading to assume that the "average VSL" implied by the range of available studies is a good measure of WTP to reduce the specific risk in the specific affected population for the policy under consideration.

The Council agrees that it is important to look at how estimated VSLs depend on characteristics of the individual (e.g., age, life expectancy), characteristics of the risk (e.g., latency, accompanying morbidity, voluntariness), and any other relevant factors. To the extent that WTP may not be a precisely proportional function of the size of the risk change, it will also be important to look more closely at the relationship between WTP estimates for different studies, concerning different specified risk changes, and to assess whether the proportionality assumption is generally tenable.

D.4. Problems with Meta-analyses

The meta-analysis in the Kochi paper, like many other meta-analyses, is premised on the assumption that there is a simple VSL relationship that is merely revealed with different degrees of bias and noise by different studies. At best, unfortunately, the

underlying construct is probably a complex VSL function. This function has many, many arguments. VSL is known to depend on the nature of the risk (severity, latency, voluntariness, etc.) and on the attributes of the individual who is considering this risk (age, gender, health status, etc.). VSL is also likely to depend upon the manner in which the demand information behind it is elicited (from self-selected employment decisions, housing choices, stated preference surveys, etc.). If only this last source of heterogeneity existed, we might be confident that techniques for pooling VSL estimates across studies would be a sensible exercise. Unfortunately, we can be fairly confident that there is fundamental heterogeneity in preferences with respect to risk, so that there is no reason, a priori, to expect that any summary statistic across studies corresponds to any single underlying “true” VSL.

The distribution of VSLs to be “averaged” in a meta-analysis is an artifact of the range of contexts (types of risks and affected populations) analyzed in the list of studies contributing to the meta-analysis. If this distribution of contexts does not correspond to the context pertinent to the environmental policy in question, then the “meta-analysis VSL” may have little to do with people’s willingness to pay the costs of this policy.

D.5. WTP and Incomplete Information

It is important to recognize two explanations for why people’s empirical decisions about mortality risk may differ from conventional theory: (1.) the individuals may be ill-informed or may make mistakes (e.g., cognitive errors), and (2) the theory may be oversimplified or wrong. It is likely that most people would like to make decisions in a way that optimizes their risk reduction spending (i.e., equal marginal spending per unit risk reduction) across various domains (e.g., housing, employment choices). However, they do not do so in practice because of information limitations and well-known errors in decision making about risk.

Some published research has made an attempt to sort out which of the factors that lead to differences between perceived risk and simple theory are simply cognitive errors (e.g., susceptibility to framing effects), and which are attributes of preferences potentially meriting normative recognition (e.g., distribution of benefits and risks of activity; such as voluntariness) (see Hammitt, 2000b).

In general, economists are inclined to defer to “consumer sovereignty” in measuring the types of tradeoffs people are willing to make. In the event of misinformation or cognitive problems, however, good policy should probably over-ride consumer errors where possible and simulate what would have been consumers’ WTP under similar conditions, but with complete and accurate information.

D.6. What to do in the near term

The Agency needs to verify that the distribution of risk reductions over which each meta-analysis has been estimated, and the context for these reductions, at least corresponds to the types of risk reductions relevant to the Clean Air Act and its amendments. The Panel continues to support meta-analyses of willingness to pay for risk reductions, but discourages the Agency from leaving the impression that it is searching for a single one-size-fits-all VSL. Instead, it should be a maintained hypothesis that heterogeneity matters. Heterogeneity should be ignored only if it can be shown to be inconsequential. The benefits from mortality (and morbidity) risk reduction attributed to a particular policy should be commensurate with the size and nature of the risk reduction and with the attributes of the affected populations.

It seems worth speculating that researchers' habit of talking in terms of conventional VSLs has much to do with the recent public relations problems concerning the "senior death discount." This different VSL for seniors was embodied in the alternative net benefits calculations associated with some recent analyses by the Agency. The public backlash to this differential seems to have been attributable almost entirely to the use of the VSL concept, which led the public to think that the issue at stake is the "value of a senior." In reality, the issue at stake is much closer to "how much money should seniors be required to pay for small risk reductions." It is essential to steer the press and the public towards the legitimacy of individual preferences and the corresponding demands (consumer sovereignty), rather than sticking with the arbitrary unit choice that expresses a marginal rate of substitution between risk changes and income as the "value of life." The word "value" is assumed by non-economists to be something intrinsic. Demand for risk reductions is not intrinsic and immutable, independent of context. It is subjective and individual, and measured differences in this demand across subpopulations and risk contexts should be honored wherever they are verifiable and based on complete information about those risks.

If WTP for small risk reductions can be shown to be approximately proportional to the size of these risk reductions over the relevant domain of the WTP function, the Panel believes it would be less inflammatory to present the marginal rate of substitution expression in terms of risk changes of a size that are pertinent to policy choices. The Panel recommends that the Agency consider converting VSL estimates into units with a less potentially misleading denominator (micromorts, millimorts, picomorts, etc.) and presenting these estimates in tandem with ordinary VSL estimates, if not in lieu of them.